

Scalability and performance improvements in Fermilab Mass Storage System.

Matt Crawford, Catalin Dumitrescu, Dmitry Litvintsev, Alexander Moibenko, Gene Oleynik

By 2009 the Fermilab Mass Storage System had encountered several challenges:

1. The required amount of data stored and accessed in both tiers of the system (dCache and Enstore) had significantly increased.
2. The number of clients accessing Mass Storage System had also increased from tens to hundreds of nodes and from hundreds to thousands of parallel requests.

To address these challenges Enstore and the SRM part of dCache were modified to scale for performance, access rates, and capacity. This work increased the amount of simultaneously processed requests in a single Enstore Library instance from about 1000 to 30000. The rates of incoming request to Enstore increased from tens to hundreds per second.

Fermilab is invested in LTO4 tape technology and we have investigated both LTO5 and Oracle T10000C to cope with the increasing needs in capacity. We have decided to adopt T10000C, mainly due to its large capacity, which allows us to scale up the existing robotic storage space by a factor 6.

This paper describes the modifications and investigations that allowed us to meet these scalability and performance challenges and provided some perspectives of Fermilab Mass Storage System.